FP7-ICT-2011.8
Contract no.: 318493

www.toposys.org

Deliverable D1.2

Progress and Activity Report on WP1

Deliverable Nature:	Report (R)
Dissemination Level: (Confidentiality)	Public (PU)
Contractual Delivery Date:	M24
Actual Delivery Date:	M24
Version:	1.0

Contents

1 Summary

In this work package, several milestones were reached. An initial implementation of the persistence of a self-map as well as the cohomology of a recurrent systems was delivered. These results were reported almost in their final form in last year's report and so are omitted in this report.

In this year we report on progress on computing the eigenvalues of a self- map as well as on a new complex construction which has applications in the entire project (including all of WP1).

The computation of all eigenvalues is based on notes and is currently a work in progress. The complex construction is based on the following paper:

- The Morse Theory of Cech and Delaunay Filtrations

Here we also highlight other pieces of work which are relevant but are still in progress.

As described in the report on WP4, the self-map algorithm does not scale well for two reasons:

1. Use of the Rips complex

2. The algorithm which explicitly computes eigenspaces

There is current work to adapt the work on the Cech and Delaunay work to self-maps (and maps in general) as well as a more efficient algorithm, which we expect to be implemented by the end of the year.

Finally, this year while we do not go into detail of the work for Hierarchies of Maps (Task 3) and Unified Theory (Task 4), we do outline the progress made. Firstly, for hierarchies of maps, i.e. diagrams which are trees, there exists a natural Heyting algebra and the recent advances in understanding on how to compute homology using variable-time set theory, one of the first examples we are examining are the cases of trees. Using the persistence lattice described in last years' report, the indecomposables of a tree were explicitly computed. We have found, however, that the variable-time set theory approach should supercede this approach (both theoretically and algorithmically) and so have not pursued it further. Once the example is computed, we do have the original approach for comparison.

In the context of unifying the theory, a key question which will be examined in the final year will be understanding different shapes which may be interesting for the self-map. In one sense we have a common language as we have a categorical notion of many of the approaches developed in this project (as highlighted in WP3). In particular, we have seen we can measure features through eigenspaces which are a form of equalizers. Dually, we can also form co-equalizers. Initial investigation has made clear that these construction contain different information. Combining these with different shapes (or perhaps restricting to maps with additional properties) will yield additional insight.

2 Finding Eigenvalues for a Self Map

Here we summarize the Kronecker theory of eigenspaces of matrix pencils, as well as some algorithm to compute the Kronecker form and (generalized) eigenvectors of a pencil, and shows some results obtained by using Kronecker form computations to compute the persistence of eigenvectors of pairs of linear maps induced in homology by simplicial maps along a filtration of simplicial complexes. In particular we go through several examples illustrating the algorithms. The goals of this work are twofold

- Compute all eigenvalues of a self map
- Understand the phenomenon of all values in a finite field becoming eigenvalues (the singular subspace)

Definition 2.1. *By the term* linear matrix pencil*, or simply* matrix pencil*, we will refer to the polynomial matrix* $\lambda B - A$*, where* $A, B \in M^{m \times n}(\mathbb{F})$*,* $\mathbb{F}$ *a field. Suppose that the equation*

$$(\lambda B - A)x(\lambda) = 0$$

possesses, for a particular value $\lambda = \lambda_0$*, a vector solution* $x(\lambda_0) \in \mathbb{F}^n$*. In this case,* λ_0 *is called an* eigenvalue *of the pencil, while* $x(\lambda_0)$ *is an* eigenvector *for this particular eigenvalue. If all values* $\lambda \in \mathbb{F}$ *are eigenvalues of the pencil, then* $x(\lambda)$ *will be a polynomial vector that solves the equation for every* λ*. In addition, suppose that we have a finite sequence* $x_1(\lambda), x_2(\lambda), \ldots, x_k(\lambda)$ *of vectors such that the system*

$$\begin{aligned} (\lambda B - A)x_1(\lambda) &= 0, \\ (\lambda B - A)x_2(\lambda) &= Bx_1(\lambda), \\ &\vdots \\ (\lambda B - A)x_k(\lambda) &= Bx_{k-1}(\lambda) \end{aligned}$$

has a solution for $\lambda = \lambda_0$*. The sequence* $x_1(\lambda_0), x_2(\lambda_0), \ldots, x_k(\lambda_0)$ *is then called a* sequence of generalized eigenvectors *for the eigenvalue* λ_0*.*

In this note we will describe what is known about the eigenvalues, eigenvectors and generalized eigenvectors of matrix pencils, using the theory of invariant polynomials of polynomial matrices (see for example [1, Chapter 6]). Let $\mathbb{F}[\lambda]$ denote the ring of polynomials in one variable over the field $\mathbb{F}$, and $P(\lambda) \in M^{m \times n}(\mathbb{F}[\lambda])$ be a polynomial matrix of dimensions $m \times n$. It is known that there exist invertible matrices $Q(\lambda) \in M^{m \times m}(\mathbb{F}[\lambda]), R(\lambda) \in M^{n \times n}(\mathbb{F}[\lambda])$ such that

$$Q(\lambda)^{-1} P(\lambda) R(\lambda) = D(\lambda) = \mathrm{diag}\{i_r(\lambda), i_{r-1}(\lambda), \ldots, i_1(\lambda), 0, \ldots, 0\}, \tag{1}$$

where the i_j's are monic polynomials with the property that $i_{j+1} | i_j$ for $j = 1, \ldots, r-1$. They are known as the *invariant polynomials* of the polynomial matrix $P(\lambda)$. In addition, $Q(\lambda)$ is a product of elementary row matrices and $R(\lambda)$, of elementary column matrices (with operations in $\mathbb{F}[\lambda]$). Such a factorization can be obtained by the classical Smith decomposition. In the case where $P(\lambda) = \lambda I - A$, where $A \in M^{n \times n}(\mathbb{F})$ is a (constant) $n \times n$ matrix and I is the identity $n \times n$ matrix over $\mathbb{F}$, there are no zeros in diagonal entries of $D(\lambda)$, and the invariant polynomials are typically said to be those of the matrix A, with i_1 commonly called the *minimal polynomial* of A, while $\prod_{j=1}^{r} i_j$ is known as the *characteristic polynomial* of A. Compare [2, Chapter 12] where the rational canonical form of a matrix $A \in M^{n \times n}(\mathbb{F})$ is computed.

The invariant polynomials can also be computed by considering minors of the polynomial matrix $P(\lambda)$, that is, determinants of square submatrices of $P(\lambda)$. Let $D_0(\lambda) \equiv 1$, and for $j = 1, \ldots, r$, $D_j(\lambda)$ be the greatest common divisor of $j \times j$ minors of $P(\lambda)$. Then

$$i_1(\lambda) = \frac{D_r(\lambda)}{D_{r-1}(\lambda)},\ i_2(\lambda) = \frac{D_{r-1}(\lambda)}{D_{r-2}(\lambda)},\ \ldots,\ i_r(\lambda) = \frac{D_1(\lambda)}{D_0(\lambda)}.$$

This can also be useful to compute the invariant polynomials of polynomial matrices in two or more variables, as we will see later.

Definition 2.2. *Let $i_1, \dots, i_r$ be the invariant polynomials of the polynomial matrix $P(\lambda)$. Write them as products of irreducible polynomials over the field $\mathbb{F}$:*

$$\begin{aligned}
i_1(\lambda) &= (p_1(\lambda))^{l_{11}}(p_2(\lambda))^{l_{12}}\dots(p_l(\lambda))^{l_{1s}}\\
i_2(\lambda) &= (p_1(\lambda))^{l_{21}}(p_2(\lambda))^{l_{22}}\dots(p_l(\lambda))^{l_{2s}}\\
&\vdots\\
i_r(\lambda) &= (p_1(\lambda))^{l_{r1}}(p_2(\lambda))^{l_{r2}}\dots(p_l(\lambda))^{l_{rs}}
\end{aligned}$$

where the exponents $l_{1k} \geqslant l_{2k} \geqslant \dots \geqslant l_{rk} \geqslant 0$, $k = 1, \dots, s$. The above irreducible factors (other than 1) are known as the elementary divisors *of the polynomial matrix $P(\lambda)$ (or, if $P(\lambda) = \lambda I - A$, equivalently of the matrix A).*

Note that unlike its invariant polynomials, the elementary divisors of a polynomial matrix $P(\lambda)$ depend on the choice of field $\mathbb{F}$. In the case where $\mathbb{F}$ is an algebraically closed field, such as $\mathbb{C}$, they will all be linear polynomials (possibly with exponent) of the form $(\lambda - \lambda_0)^l$. This allows one to formally define the Jordan normal form (for square matrices) over such a field, with the exponents of the elementary divisors being the dimension of the Jordan blocks. The same is true for matrix pencils, where the Kronecker canonical form, being a generalization of the Jordan form, will only be guaranteed to exist when working on an algebraically closed field. In the sequel, it is assumed that we are working over $\mathbb{C}$, unless otherwise stated.

Definition 2.3. *Let $\lambda B - A$ be a matrix pencil in $\mathbb{C}$. If it is square ($B, A \in M^{n\times n}(\mathbb{C})$) and its determinant does not vanish identically, it is said to be* regular*. If it is non-square, or square but its determinant is identically 0, it is said to be* singular*.*

Let $\lambda B_1 - A_1$, $\lambda B_2 - A_2$ be two $m \times n$ pencils. If there exist invertible matrices $P \in M^{m\times m}(\mathbb{C})$, $Q \in M^{n\times n}(\mathbb{C})$ such that $P^{-1}(\lambda B_1 - A_1)Q = \lambda B_2 - A_2$, then the pencils are said to be similar*. (In [1], this concept of similarity is referred to as* strict equivalence *given that the transformation matrices are constant.)*

Now, if the matrices $A_1, A_2 \in M^{n\times n}(\mathbb{C})$ have the same elementary divisors, then they have the same Jordan normal form J, and therefore there exist matrices T_1 and T_2, whose columns are respectively generalized eigenvectors of A_1 and A_2, such that $J = T_1^{-1}A_1T_1 = T_2^{-1}A_2T_2$. Therefore, the pencils $\lambda I - A_1$ and $\lambda I - A_2$ are both similar to $\lambda I - J$, and therefore to each other. This, however, is not true even in the case of regular pencils: two regular pencils $\lambda B_1 - A_1$ and $\lambda B_2 - A_2$ can have the same elementary divisors without being similar.

Example 2.4. *[1, Chapter 12]* Let

$$\lambda B_1 - A_1 = \lambda \begin{bmatrix} 1 & 1 & 2 \\ 1 & 1 & 2 \\ 1 & 1 & 3 \end{bmatrix} - \begin{bmatrix} -2 & -1 & -3 \\ -3 & -2 & -5 \\ -3 & -2 & -6 \end{bmatrix},$$
$$\lambda B_2 - A_2 = \lambda \begin{bmatrix} 1 & 1 & 1 \\ 1 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix} - \begin{bmatrix} -2 & -1 & -1 \\ -1 & -2 & -1 \\ -1 & -1 & -1 \end{bmatrix}.$$

It is easy to show that the two previous pencils each only have a single elementary divisor, that is, $\lambda + 1$. However, they are clearly not similar, since B_1 has rank 2, while B_2 only has rank 1.

In order to explain this phenomenon, we will compute the elementary divisors of the *homogeneous* pencils $\lambda B_1 - \mu A_1$ and $\lambda B_2 - \mu A_2$.

$$\lambda B_1 - \mu A_1 = \begin{bmatrix} \lambda + 2\mu & \lambda + \mu & 2\lambda + 3\mu \\ \lambda + 3\mu & \lambda + 2\mu & 2\lambda + 5\mu \\ \lambda + 3\mu & \lambda + 2\mu & 3\lambda + 6\mu \end{bmatrix},$$

$$\lambda B_2 - \mu A_2 = \begin{bmatrix} \lambda + 2\mu & \lambda + \mu & \lambda + \mu \\ \lambda + \mu & \lambda + 2\mu & \lambda + \mu \\ \lambda + \mu & \lambda + \mu & \lambda + \mu \end{bmatrix}$$

In the case of $\lambda B_1 - \mu A_1$, we compute $D_1(\lambda, \mu) = 1$, $D_2(\lambda, \mu) = 1$, and $D_3(\lambda, \mu) = \mu^2(\lambda + \mu)$. In other words, $i_1(\lambda, \mu) = \mu^2(\lambda + \mu)$ is the only nontrivial invariant polynomial of $\lambda B_1 - \mu A_1$, and so we have as elementary divisors μ^2 and $\lambda + \mu$. For $\lambda B_2 - \mu A_2$, $D_1(\lambda, \mu) = 1$, $D_2(\lambda, \mu) = \mu$, and $D_3(\lambda, \mu) = \mu^2(\lambda + \mu)$. There are two nontrivial invariant polynomials: $i_1(\lambda, \mu) = \mu(\lambda + \mu)$ and $i_2(\lambda, \mu) = \mu$, and so the elementary divisors are μ, $-\mu$, and $\lambda + \mu$.

This suggests to us the existence of *infinite elementary divisors* created by the matrix holding the λ parameter being singular. By an arbitrarily small deformation, we could create a pencil with eigenvalues arbitrarily large in modulus, and so it makes sense to say that the pencil has "infinity" as an eigenvalue. In the previous example, $\lambda B_1 - A_1$ has one infinite elementary divisor of order 2, while $\lambda B_2 - A_2$ has two infinite elementary divisors of order 1. As for the "ordinary" elementary divisors ($\lambda + 1$, in the previous example), they will be referred to as *finite elementary divisors*.

In other words, regular pencils $\lambda B - A$ have both a finite and an infinite eigenstructure. We can show (see [1, Chapter 12]) that the pencil $\lambda B - A$ is strictly equivalent to

$$\operatorname{diag}\{\lambda N - I, \lambda I - J\}$$

where N is a nilpotent Jordan matrix, and where J is in Jordan canonical form. The pencil $\lambda N - I$ is associated with the infinite eigenstructure of $\lambda B - A$, while $\lambda I - J$ contains its finite eigenstructure. In other words, N contains Jordan blocks (associated with the eigenvalue 0) corresponding to the infinite elementary divisors, while J contains Jordan blocks (associated with the finite eigenvalues of the pencil) corresponding to the finite elementary divisors.

For example, in Example 2.4, the pencil $\lambda B_1 - A_1$ will be similar to

$$\begin{bmatrix} -1 & \lambda & \\ & -1 & \\ & & \lambda + 1 \end{bmatrix}$$

where we see the 2×2 block associated with the infinite elementary divisor of order 2, while $\lambda B_2 - A_2$ will instead be similar to

$$\begin{bmatrix} -1 & & \\ & -1 & \\ & & \lambda + 1 \end{bmatrix}$$

where we instead have two 1×1 blocks associated with both of the infinite elementary divisors of order 1.

Note that if λ_0 is a finite, nonzero eigenvalue of $\lambda B - A$, then $1/\lambda_0$ is an eigenvalue of $\lambda A - B$. In addition, if λ^r is an elementary divisor of $\lambda B - A$ (i.e. 0 is an eigenvalue of $\lambda B - A$ associated with a $r \times r$ Jordan block), then $\lambda A - B$ has a r-th order infinite elementary divisor, and vice-versa.

Now let us go back to singular pencils $\lambda B - A$ of dimension $m \times n$. If the set of columns of the pencil is linearly dependent, then there exists a polynomial eigenvector, that is, a solution to $(\lambda B - A)x(\lambda) = 0$ of the form

$$x(\lambda) = x_0 + \lambda x_1 + \lambda^2 x_2 + \ldots + \lambda^\epsilon x_\epsilon,\ \epsilon \geqslant 0. \tag{2}$$

By re-writing this system as

$$\begin{aligned} Ax_0 &= 0 \\ Ax_1 &= Bx_0 \\ Ax_2 &= Bx_1 \\ &\vdots \\ Ax_\epsilon &= Bx_{\epsilon-1} \end{aligned}$$

we see that finding a polynomial solution $x(\lambda)$ is equivalent to finding a solution to the system

$$T_\epsilon(A, B)x = 0$$

where $x = [x_0\, x_1\, x_2\, \ldots\, x_\epsilon]^T$ and

$$T_\epsilon(A, B) = \begin{bmatrix} A & & & & \\ -B & A & & & \\ & -B & A & & \\ & & \ddots & \ddots & \\ & & & -B & A \\ & & & & -B \end{bmatrix} \in M^{(\epsilon+2)m \times (\epsilon+1)n}(\mathbb{C})$$

is called the ϵ*-th Toeplitz matrix* of (A, B). As a matter of fact, the lowest non-negative integer ϵ such that $(\lambda B - A)x(\lambda) = 0$ possesses a solution of degree ϵ can be computed as the lowest ϵ such that $\mathsf{rank}\, T_\epsilon(A, B)$ is strictly smaller than $(\epsilon + 1)n$. Other algorithms to compute the Kronecker canonical form of a pencil (or a similar pencil in which all the information included in the Kronecker form is easily accessible) also provide this ϵ, as we shall see.

Theorem 2.5. *If the pencil $\lambda B - A$ possesses a polynomial solution (2) of minimal degree $\epsilon > 0$, then it is similar to the pencil*

$$\begin{bmatrix} L_\epsilon & 0 \\ 0 & \lambda\hat{B} - \hat{A} \end{bmatrix}$$

where

$$L_\epsilon = \begin{bmatrix} \lambda & -1 & & \\ & \ddots & \ddots & \\ & & \lambda & -1 \end{bmatrix}$$

is a bidiagonal pencil of dimension $\epsilon \times (\epsilon + 1)$ and $\lambda\hat{B} - \hat{A}$ has no solution analogous to (2) of degree less than ϵ. This theorem can be found as Theorem 4 in [1, Chapter 12]. Note that we allow blocks of dimension 0×1 (meaning zero columns in the pencil we are constructing), which is slightly different from how it is done in this book.

The previous theorem can be applied repeatedly to extract from the pencil $\lambda B - A$ all the linear dependence present in its set of columns. Therefore, by a similarity transformation applied on $\lambda B - A$, we can obtain

$$\begin{bmatrix} L_{\epsilon_1} & & & & \\ & L_{\epsilon_2} & & & \\ & & \ddots & & \\ & & & L_{\epsilon_p} & \\ & & & & \lambda B_p - A_p \end{bmatrix}$$

where $0 \leqslant \epsilon_1 \leqslant \epsilon_2 \leqslant \ldots \leqslant \epsilon_p$ are known as the *Kronecker column indices*, and the pencil $\lambda B_p - A_p$ has no linear dependence on its columns. If it has linear dependence on its rows, we can apply the theorem to the transposed pencil $\lambda B_p^T - A_p^T$, obtaining *Kronecker row indices* $0 \leqslant \eta_1 \leqslant \eta_2 \leqslant \ldots \leqslant \eta_q$. Therefore, the pencil $\lambda B - A$ is similar to

$$\begin{bmatrix} L_{\epsilon_1} & & & & & & & & \\ & L_{\epsilon_2} & & & & & & & \\ & & \ddots & & & & & & \\ & & & L_{\epsilon_p} & & & & & \\ & & & & L_{\eta_1}^T & & & & \\ & & & & & L_{\eta_2}^T & & & \\ & & & & & & \ddots & & \\ & & & & & & & L_{\eta_q}^T & \\ & & & & & & & & \lambda B_0 - A_0 \end{bmatrix}$$

where there is no linear dependence on the rows or columns of $\lambda B_0 - A_0$, which is therefore a regular pencil (and hence square).

The Kronecker blocks L_ϵ (resp. L_η) encode the singularity on the columns (resp. rows) of $\lambda B - A$, which we can see by noticing that

$$L_\epsilon \begin{bmatrix} 1 \\ \lambda \\ \lambda^2 \\ \vdots \\ \lambda^\epsilon \end{bmatrix} = \begin{bmatrix} \lambda & -1 & & \\ & \ddots & \ddots & \\ & & \lambda & -1 \end{bmatrix} \begin{bmatrix} 1 \\ \lambda \\ \lambda^2 \\ \vdots \\ \lambda^\epsilon \end{bmatrix} = 0$$

for all field values λ. The same is true when pre-multiplying the blocks L_η^T by $[1\,\lambda\,\lambda^2\,\cdots\,\lambda^\eta]$.

This now allows us to state that every $m \times n$ pencil $\lambda B - A$ can be put by a similarity transformation into the *Kronecker canonical form*:

$$\mathrm{diag}\{L_{\epsilon_1}, \ldots, L_{\epsilon_p}, L_{\eta_1}^T, \ldots, L_{\eta_q}^T, \lambda N - I, \lambda I - J\}$$

where the ϵ_i's are the column Kronecker indices, the η_j's the row Kronecker indices, N a nilpotent Jordan matrix ($\lambda N - I$ encodes the infinite eigenstructure of the pencil) and J in Jordan form ($\lambda I - J$ encodes the finite eigenstructure of the pencil).

Example 2.6. Consider the matrix pencil $\lambda B - A$, where

$$B = \begin{bmatrix} 6 & -5 & 0 & 1 & -1 & 1 & 1 \\ 2 & -2 & 0 & 0 & 0 & 0 & 0 \\ 0 & -3 & -2 & 1 & -1 & -1 & 0 \\ 1 & 8 & 6 & -1 & 0 & 0 & 0 \\ 0 & -6 & -4 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & -1 & 1 & 1 \end{bmatrix}, \; A = \begin{bmatrix} -5 & 5 & -1 & 0 & -4 & -6 & -1 \\ -2 & 2 & 0 & 0 & -2 & -2 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ -1 & 1 & 0 & 0 & -1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & -2 & 0 \\ 1 & 0 & -1 & 1 & 1 & 1 & 0 \end{bmatrix}.$$

By a similarity transformation we find that the Kronecker canonical form of this pencil is

$$\begin{bmatrix} \lambda & -1 & & & & & \\ & & -1 & \lambda & & & \\ & & 0 & -1 & & & \\ & & & & \lambda & & \\ & & & & & \lambda+1 & -1 \\ & & & & & 0 & \lambda+1 \end{bmatrix}.$$

In other words, the column Kronecker index $\epsilon_1 = 1$ corresponds to a two-dimensional singular structure, the submatrix of the form $\lambda N - I$ corresponds to an infinite elementary divisor of order 2, and the finite structure shows the elementary divisors λ and $(\lambda + 1)^2$.

This being said, we require an algorithm to put a pencil $\lambda B - A$ into Kronecker form. Van Dooren [3] describes an algorithm using unitary transformations to effect row and column reductions on submatrices of $\lambda B - A$ ("compressions"), putting it after a finite l number of steps into block lower triangular form, precisely,

$$\begin{array}{c} \left[\begin{array}{c|c|c|c|c} \lambda B_{l+1,l+1} - A_{l+1,l+1} & 0 & \cdots & 0 & 0 \\ \hline \lambda B_{l+1,l} - A_{l+1,l} & -A_{l,l} & \cdots & 0 & 0 \\ \hline \vdots & \ddots & \ddots & \vdots & \vdots \\ \hline \lambda B_{l+1,2} - A_{l+1,2} & \lambda B_{l,2} - A_{l,2} & \cdots & -A_{2,2} & 0 \\ \hline \lambda B_{l+1,1} - A_{l+1,1} & \lambda B_{l,1} - A_{l,1} & \cdots & \lambda B_{2,1} - A_{2,1} & -A_{1,1} \end{array}\right] \begin{array}{l} m_{l+1} \\ r_l \\ \vdots \\ v_2 \\ v_1 \end{array} \\ \begin{array}{ccccc} n_{l+1} & s_l & \cdots & s_2 & s_1 \end{array} \end{array} \tag{3}$$

where the $(l + 1) \times (l + 1)$ blocks are delimited with vertical and horizontal lines, $B_{l+1,l+1}$ has full column rank, the $A_{i,i}$'s have full row rank r_i for $i = 1, \ldots, l$, and the $B_{i,i-1}$'s have full column rank s_i for $i = 2, \ldots, l$. From the row and column ranks r_i and s_i, we then compute (putting $s_{l+1} = 0$)

$$s_i - r_i = e_i \geqslant 0 \text{ for } i = 1, \ldots, l;$$
$$r_i - s_{i+1} = d_i \geqslant 0 \text{ for } i = 1, \ldots, l.$$

The indices d_i and e_i fully determine the infinite elementary divisors and the column Kronecker indices, respectively. More precisely, they tell us that the pencil $\lambda B - A$ has d_i infinite elementary divisors of degree i, $i = 1, \ldots, l$, and e_i column Kronecker blocks L_{i-1} of size $(i-1)\times i$, $i = 1, \ldots, l$. A similar but "dual" algorithm may then be used on $\lambda B_{l+1,l+1} - A_{l+1,l+1}$ to extract the row Kronecker indices, but since row Kronecker blocks do not possess eigenvectors, this step may not be necessary for our purposes. The pencil $\lambda B_{l+1,l+1} - A_{l+1,l+1}$ additionally contains the finite structure of the original pencil.

Through similarity transformations, we can further reduce $\lambda B - A$ to the form $\mathrm{diag}\{\lambda B_{l+1,l+1} - A_{l+1,l+1}, \lambda B_n - A_n\}$, where

$$\lambda B_n - A_n = \begin{bmatrix} -J_l & & & \\ \lambda K_{l-1} & -J_{l-1} & & \\ & \ddots & \ddots & \\ & & \lambda K_1 & -J_1 \end{bmatrix}$$

and

$$J_i = \left[\begin{array}{c|c} I_{r_i} & 0 \end{array}\right], \quad K_i = \left[\begin{array}{c} I_{s_{i+1}} \\ \hline 0 \end{array}\right]$$

with J_i having s_i columns and K_i having r_i rows. At this point, we have invertible matrices Q and R such that $Q^{-1}(\lambda B - A)R = \mathrm{diag}\{\lambda B_{l+1,l+1} - A_{l+1,l+1}, \lambda B_n - A_n\}$. This is as far as we can get only with similarity transformations, but by applying row and column operations with polynomial coefficients (normal – non-strict – *equivalence* transformations, in the language of [1]), we can obtain polynomial eigenvectors for the pencil $\lambda B - A$. The following example should make this process clearer.

Example 2.7. Let us go back to the pencil of Example 2.6, where the base field is considered to be $\mathbb{Z}_{19}$. We have used this field as it is finite, and therefore easier for computational purposes, but also decently large and with some potential for different eigenvalues; of course other fields would be possible. It's understood that in this field $-n$ means the element $19 - n$. We have used the row and column reduction methods of the CAPD library for the row and column compressions required in the algorithm of [3].

Step 1: $r_1 = 2, s_1 = 2$,

$$B \sim \left[\begin{array}{ccccc|cc} -1 & 7 & 4 & 0 & -1 & 0 & 0 \\ 0 & -6 & 0 & 0 & 0 & 0 & 0 \\ 1 & -6 & -4 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 & 0 & 0 \\ \hline -1 & 7 & -9 & 1 & 0 & 0 & 0 \\ 8 & 8 & 6 & -7 & 0 & 0 & 0 \end{array}\right], A \sim \left[\begin{array}{ccccc|cc} 0 & -1 & 0 & 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 6 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & -2 & 0 & 0 & -2 & 0 & 0 \\ \hline 0 & -1 & 0 & 0 & -1 & 2 & 0 \\ 0 & 2 & 0 & -1 & 2 & 0 & 7 \end{array}\right];$$

Step 2: $r_2 = 1, s_2 = 2$,

$$B \sim \left[\begin{array}{ccc|cc|cc} 0 & 2 & 0 & 0 & 0 & 0 & 0 \\ 4 & 0 & 0 & 0 & 0 & 0 & 0 \\ -4 & 1 & -1 & 0 & 0 & 0 & 0 \\ \hline 0 & -6 & 0 & 0 & 0 & 0 & 0 \\ \hline -9 & -1 & -7 & 1 & -8 & 0 & 0 \\ 6 & 7 & -8 & -7 & 0 & 0 & 0 \end{array}\right], A \sim \left[\begin{array}{ccc|cc|cc} 0 & -2 & -2 & 0 & 0 & 0 & 0 \\ 0 & -1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 6 & 0 & 1 & 0 & 0 \\ \hline 0 & -1 & -1 & 0 & 0 & 2 & 0 \\ 0 & 2 & 2 & -1 & 0 & 0 & 7 \end{array}\right];$$

Step 3: $r_3 = 0, s_3 = 0$,

$$B \sim \left[\begin{array}{ccc|cc|cc} 2 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 4 & 0 & 0 & 0 & 0 & 0 \\ 1 & -4 & -1 & 0 & 0 & 0 & 0 \\ \hline -6 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline -1 & -9 & -7 & 1 & -8 & 0 & 0 \\ 7 & 6 & -8 & -7 & 0 & 0 & 0 \end{array}\right], A \sim \left[\begin{array}{ccc|cc|cc} -2 & 0 & -2 & 0 & 0 & 0 & 0 \\ -1 & 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 6 & 0 & 1 & 0 & 0 \\ \hline -1 & 0 & -1 & 0 & 0 & 2 & 0 \\ 2 & 0 & 2 & -1 & 0 & 0 & 7 \end{array}\right].$$

The algorithm's stopping criterion is reached. We find that $e_2 = 1$, so there is one 1×2 column Kronecker block, and $d_2 = 1$, so there is one infinite elementary divisor of order 2. The pencil $\lambda B_{3,3} - A_{3,3}$ is as well of size 3×3, so since it is square, we additionally know that there is no row Kronecker block. We apply the second part of the algorithm, to zero out subdiagonal blocks:

Step 1:

$$B \sim \left[\begin{array}{ccc|cc|cc} 1 & -4 & -1 & 0 & 0 & 0 & 0 \\ 0 & 4 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline 8 & -9 & -7 & -8 & 1 & 0 & 0 \\ 7 & 6 & -8 & 0 & -7 & 0 & 0 \end{array}\right], \; A \sim \left[\begin{array}{ccc|cc|cc} 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ -1 & 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & -2 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ \hline -1 & 0 & -1 & 0 & 0 & 2 & 0 \\ 2 & 0 & 2 & 0 & -1 & 0 & 7 \end{array}\right];$$

Step 2:

$$B \sim \left[\begin{array}{ccc|cc|cc} 1 & -4 & -1 & 0 & 0 & 0 & 0 \\ 0 & 4 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & -8 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & -7 & 0 & 0 \end{array}\right], \; A \sim \left[\begin{array}{ccc|cc|cc} 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ -1 & 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & -2 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 7 \end{array}\right].$$

Written as a block pencil,

$$\lambda B - A \sim \left[\begin{array}{ccc|cc|cc} \lambda & -4\lambda & -1\lambda - 1 & & & & \\ 1 & 4\lambda & 1 & & & & \\ 0 & 0 & 2\lambda + 2 & & & & \\ \hline & & & -1 & 0 & & \\ \cline{4-7} & & & -8\lambda & \lambda & -2 & 0 \\ & & & 0 & -7\lambda & 0 & -7 \end{array}\right]$$

with the column transformation matrix being

$$R = \begin{bmatrix} 7 & 1 & 6 & 5 & -2 & -2 & -1 \\ 6 & 1 & 6 & 5 & -2 & -2 & -1 \\ -7 & 0 & -6 & 0 & 8 & 3 & -8 \\ 5 & -1 & 6 & -4 & -9 & 0 & 0 \\ 0 & 0 & 7 & 0 & 0 & -1 & 0 \\ 0 & 0 & -6 & 0 & 0 & 1 & 0 \\ 8 & 0 & 1 & 0 & -8 & 0 & 1 \end{bmatrix}.$$

Using the rightmost block to zero out column 5 of the pencil, and applying the same transformation to R, we obtain the polynomial eigenvector $x_1(\lambda) = [-2, -2, 8, -9, 9\lambda, -9\lambda, -8 - \lambda]^T$ associated with the Kronecker block of index 1. Using Smith diagonalization on the top leftmost block (where the finite structure lies) and also applying the transformation to R, we obtain a single invariant polynomial, $\lambda^3 + 2\lambda^2 + \lambda$, and therefore two elementary divisors in $\mathbb{Z}_{19}$: λ and $(\lambda+1)^2$; with another polynomial eigenvector $x_2(\lambda) = [-4\lambda^2 + 6\lambda + 1, -5\lambda + 1, 4\lambda^2 - 6\lambda, 4\lambda^2 + 3\lambda - 1, 9\lambda^2 + 9\lambda, -5\lambda^2 - 5\lambda, -9\lambda^2 - 3\lambda]^T$. For every $\lambda \in \mathbb{Z}_{19}$, $x_1(\lambda)$ is an eigenvector for λ, while for $\lambda = 0$, $x_2(0) = [1, 1, 0, -1, 0, 0, 0]^T$ is another eigenvector, and for $\lambda = -1$, $x_2(-1) = [-9, 6, -9, 0, 0, 0, -6]^T$ is another.

At this point, we wonder how to compute generalized eigenvectors for a pencil. Indeed, for example, the pencil of Example 2.6 has $(\lambda+1)^2$ as elementary divisor, to which is associated a two-dimensional generalized eigenspace for eigenvalue -1, and a column Kronecker block of index 1, to which is associated as well a two-dimensional generalized eigenspace, for every field value. Suppose that $x(\lambda)$ is a polynomial eigenvector associated with a column Kronecker block. In other words,

$$(\lambda B - A)\, x(\lambda) = 0 \tag{4}$$

and this equation is valid for every $\lambda \in \mathbb{F}$. Consider formal differentiation over the ring $\mathbb{F}[\lambda]$. This is a ring homomorphism denoted $\frac{d}{d\lambda}$ (we can also use prime notation) with the property that

$$\frac{d}{d\lambda}\lambda^k = k\lambda^{k-1} \text{ for } k \in \mathbb{N}, \text{ and } \frac{d}{d\lambda}c = 0, c \in \mathbb{F}.$$

Formal differentiation verifies the chain rule, and so we can apply it repeatedly to both sides of Equation (4):

$$\begin{aligned}(\lambda B - A)x'(\lambda) &= -Bx(\lambda),\\ (\lambda B - A)x''(\lambda) &= -2Bx'(\lambda),\\ &\vdots\\ (\lambda B - A)x^{(k)}(\lambda) &= -kBx^{(k-1)}(\lambda).\end{aligned}$$

If the degree of x is k and the characteristic of $\mathbb{F}$ is strictly greater than k, we therefore get a sequence of $k+1$ generalized eigenvectors for every $\lambda \in \mathbb{F}$. If now $x(\lambda)$ is a polynomial eigenvector associated with an invariant polynomial $i(\lambda)$ of $\lambda B - A$, then we know that $i(\lambda) = (\lambda - \lambda_0)^{k+1} r(\lambda)$, where λ_0 is a root of i (an eigenvalue) of multiplicity $k+1$, $k \geqslant 0$, and $r(\lambda_0) \neq 0$. Then, by Equation 1, if $e \in \mathbb{F}^m$ is a vector with 1 in the appropriate position and zeros elsewhere,

$$(\lambda B - A)x(\lambda) = (\lambda - \lambda_0)^{k+1} r(\lambda) Q(\lambda) e.$$

The right-hand side of this equation equals 0 when $\lambda = \lambda_0$, and therefore $x(\lambda_0)$ is an eigenvector for the pair for the eigenvalue λ_0. In addition, by applying formal differentiation i times to both sides of the equation, for $i \leqslant k$, we obtain

$$\begin{aligned}(\lambda B - A)x^{(i)}(\lambda) = &- iBx^{(i-1)}(\lambda)\\ &+ \sum_{j=0}^{i} \binom{i}{j} \frac{(k+1)!}{(k+1+i-j)!} (\lambda - \lambda_0)^{k+1-j} \frac{d^{i-j}}{d\lambda^{i-j}} (r(\lambda)Q(\lambda))e,\end{aligned}$$

and the sum on the right-hand side equals 0 when $\lambda = \lambda_0$. We therefore once again obtain a sequence of $k+1$ generalized eigenvectors for λ_0.

If we return to Example 2.6, the polynomial eigenvector

$$x_1(\lambda) = [-2, -2, 8, -9, 9\lambda, -9\lambda, -8-\lambda]^T$$

associated with a Kronecker block of index 1 yields a generalized eigenvector $x_1'(\lambda) = [0, 0, 0, 0, 9, -9, -1]^T$. (We may neglect the -1 coefficient since a multiple of a generalized eigenvector is also a generalized eigenvector.) As for the polynomial eigenvector $x_2(\lambda) = [-4\lambda^2 + 6\lambda + 1, -5\lambda + 1, 4\lambda^2 - 6\lambda, 4\lambda^2 +$

$3\lambda - 1, 9\lambda^2 + 9\lambda, -5\lambda^2 - 5\lambda, -9\lambda^2 - 3\lambda]^T$, associated with the eigenvalue $\lambda = -1$ of multiplicity 2, a generalized eigenvector will be $x_2'(-1) = [-5, -5, 5, -5, -9, 5, -4]^T$.

We have applied the above method to select eigenvalues and identify eigenvectors to the algorithm described in [4], which computes the persistence of eigenspaces of pairs of linear maps in homology as a means to reconstruct the dynamics of a sampled map. The program has been run on a map on a cloud of points in $S^1 \subset \mathbb{C}$, representing the map $z \mapsto z^2$, where the samples are taken on the unit circle and then subjected to several levels of Gaussian noise with standard deviation from $\sigma = 0$ to 0.30. It is expected that we should find in homology H_1 (computed over the field $\mathbb{Z}_{19}$) an eigenvector of long persistence for eigenvalue $\lambda = 2$, but that stronger noise may make it harder to distinguish. Figure 1 shows the persistence barcodes for the eigenvector associated with $\lambda = 2$ along a filtration of complexes indexed with parameter value ϵ, colour-coded in the following way: where the eigenvector is associated with a finite elementary divisor of the form $(\lambda - 2)^r$; that is, where 2 is a finite eigenvalue of the pair of maps in homology between two successive ϵ values, the bar is drawn blue. When however the eigenvector is associated with a column Kronecker block in the canonical form of this pair of maps – and in other words actually becomes an eigenvector for every field value, not only $\lambda = 2$ – the bar is drawn red. It can be seen that as the noise level is increased, the persistence of this eigenvector tends to become shorter, being born later and dying earlier, and additionally the eigenvector becomes "degenerate" (associated with the singular structure of the pencil) for a longer term. It should be noted that in our current numerical experiments, we have witnessed a persistent eigenvector formerly associated with a finite eigenvalue become degenerate, but not yet the inverse (an eigenvector associated with the singular structure becoming associated with a unique finite eigenvalue). We however do not know any reason why this should not also be possible.

Note however that defining persistence for degenerate eigenvectors is fraught with hazard, since it is not even known yet if such a concept is theoretically possible, given that it is unknown if polynomial eigenvectors for the singular structure form a vector space. Nevertheless, when the polynomial eigenvector is evaluated at a particular field value, persistence can be computed, which is what we have done here.

3 Morse Theory Cech and Delaunay Complexes

A fundamental task in topological data analysis is to turn the data into a topological space, or a filtration of spaces. Assuming the data consists of points in $\mathbb{R}^n$, we have several possible ways to construct a filtration using a non-negative scale parameter r. We can connect $p + 1$ data points by the spanned p-simplex iff the closed balls of radius r centered at the points have a non-empty common intersection. This gives the *Čech complex* for radius r, which is known to have the same homotopy type as the union of balls [15, Chapter III]. We can intersect each ball with the Voronoi cell of its point and then take the nerve of these intersections. This gives the *alpha* or *Delaunay complex* for radius r, which embeds in $\mathbb{R}^n$ and also has the homotopy type of the union of balls [13]. Alternatively, we can select all simplices in the Čech complex that belong to the Delaunay triangulation. This gives the *Delaunay-Čech complex* for radius r, which also embeds in $\mathbb{R}^n$ but is known to collapse to the Delaunay complex only in $\mathbb{R}^2$ [10]. These results generalize this relation to arbitrary dimension.

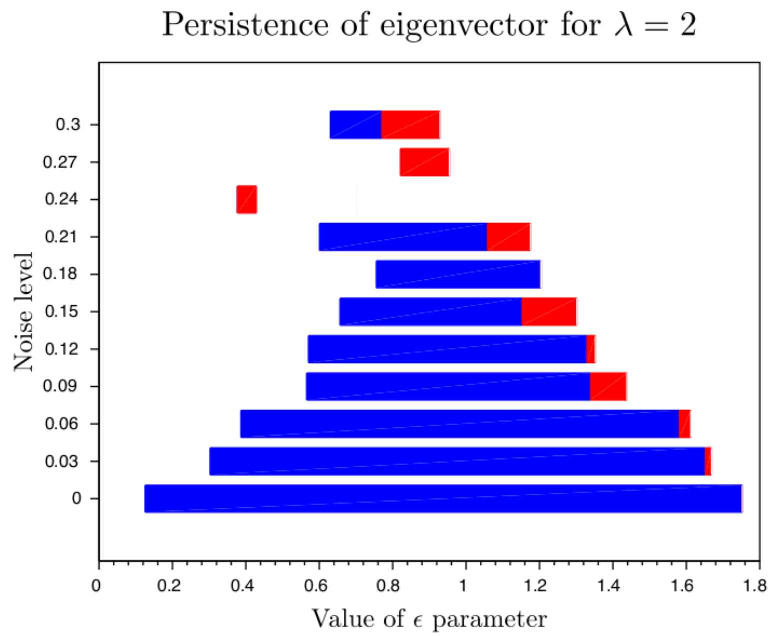

Figure 1: Persistence of the eigenvector associated with $\lambda = 2$ for several noise levels of a cloud of sample points on S^1, subject to the map $z \mapsto z^2$.

Data analysis in low dimensions usually requires the reconstruction of the shape beyond the mere topological characterization. An example is the *wrap complex*, which is used commercially for this purpose [14]. We extend the original 3-dimensional notion to $\mathbb{R}^n$ and introduce a dependence on the radius parameter, showing that the wrap complex is simple-homotopy equivalent to the Delaunay and the Delaunay-Čech complexes. Among the four, the Delaunay complex has been defined for points with weights. We generalize the other three complexes to the weighted case in a way that preserve the homotopy and simple-homotopy equivalences between them.

The main technical ingredients to proving the relations between the complexes are generalized discrete Morse functions; see [18] for an introduction to discrete Morse theory and [19] for the generalization that allows for intervals larger than pairs in the discrete vector field. Our constructions are elementary, using the radii of smallest empty circumspheres and smallest enclosing spheres to define the generalized discrete Morse functions, which are then used to prove the collapsibility results.

3.1 Background

All complexes are simplicial, either concrete geometric or abstract. Here we introduce three of the four types of complexes, and we give quick reviews of discrete Morse theory and its generalization.

To apply standard concepts from point set topology, we associate with a (geometric) simplicial complex K in $\mathbb{R}^n$ its *underlying space*, $|K|$, which is the union of the simplices of K endowed with the subspace topology inherited from $\mathbb{R}^n$. Two topological spaces $\mathbb{X}$ and $\mathbb{Y}$ are *homeomorphic* if there is a continuous bijection $h\colon \mathbb{X} \to \mathbb{Y}$ whose inverse is also continuous. A weaker but often more useful notion is the following. The two spaces have the same *homotopy type* if there are continuous maps $f\colon \mathbb{X} \to \mathbb{Y}$ and $g\colon \mathbb{Y} \to \mathbb{X}$ such that $g \circ f$ is homotopic to the identity on $\mathbb{X}$ and $f \circ g$ is homotopic to the identity of $\mathbb{Y}$. Here, two maps are said to be *homotopic* if there is a continuous deformation from one map to the other. To apply this concept in the abstract setting, we first construct geometric realizations of all abstract simplicial complexes, which is done without mention. A useful tool in this context is the Nerve Theorem, which asserts that under certain conditions, the nerve of a collection of sets has the same homotopy type as the union of the sets [6, 23]. It applies, for example, if the collection is finite and the sets are convex.

Proximity complexes Let X be a finite set of points in $\mathbb{R}^n$. For $r \geqslant 0$, let $B_r(x) = \{y \in \mathbb{R}^n \mid \|x - y\| \leqslant r\}$ be the closed ball of radius r centered at $x \in X$. The *Čech complex* for radius r,

$$C_r(X) = \{Q \subseteq X \mid \bigcap_{x \in Q} B_r(x) \neq \emptyset\},$$

is isomorphic to the nerve of the collection of closed balls. For $r \leqslant s$, we have $C_r(X) \subseteq C_s(X)$. For sufficiently large radius, the Čech complex is the full (abstract) simplex spanned by X, which we denote as $\Delta(X) = 2^X \backslash \{\emptyset\}$. Using the Euclidean metric in $\mathbb{R}^n$, we define the *Voronoi cell* of $x \in X$ as the set of points $a \in \mathbb{R}^n$ such that $\|x - a\| \leqslant \|y - a\|$ for all $y \in X$. Every Voronoi cell is a convex polyhedron, any two such polyhedra intersect at most along shared boundaries, and together the Voronoi cells cover $\mathbb{R}^n$. Letting $\mathrm{Vor}\,(x)$ be the Voronoi cell of $x \in X$, we write $\mathrm{Vor}\,(Q) = \bigcap_{x \in Q} \mathrm{Vor}\,(x)$ for the common intersection of Voronoi cells, for every $Q \subseteq X$. The *Delaunay triangulation*, denoted by $D(X)$, is isomorphic to the nerve of the collection of Voronoi cells: it consists of all simplices $Q \subseteq X$ with $\mathrm{Vor}\,(Q) \neq \emptyset$. A mild general position assumption is required to ensure that taking the convex hulls of the abstract simplices in $D(X)$ gives a simplicial complex in $\mathbb{R}^n$.

There are two natural ways to combine the two constructions. For the first, we intersect $B_r(x)$ with $\mathrm{Vor}\,(x)$ and we construct a complex isomorphic to the nerve of the collection of such sets, and for the second, we restrict the Čech complex to the Delaunay triangulation:

$$D_r(X) = \{Q \subseteq X \mid \bigcap_{x \in Q} [B_r(x) \cap \mathrm{Vor}\,(x)] \neq \emptyset\},$$

$$DC_r(X) = \{Q \in D(X) \mid \bigcap_{x \in Q} B_r(x) \neq \emptyset\}.$$

In the literature, $D_r(X)$ is known as the *alpha complex*, for $\alpha = r$; see e.g. [16]. We prefer to call it the *Delaunay complex* for radius r, because this naming convention emphasizes the fact that the Delaunay complexes provide the canonical filtration of the Delaunay triangulation, as explained in 3.2. The second complex, $DC_r(X)$, is referred to as the *Delaunay-Čech complex* for radius r. For growing radius, both complexes give rise to a filtration of the Delaunay triangulation. While the two complexes are similar, they are not necessarily the same; see 2. Instead of equality, we have $D_r(X) \subseteq DC_r(X)$, for all r. To see this, we just note that if the sets $B_r(x) \cap \mathrm{Vor}\,(x)$ have a non-empty common intersection, then the sets $B_r(x)$ have a non-empty common intersection and so do the sets $\mathrm{Vor}\,(x)$.

Discrete Morse theory Let K be a simplicial complex, geometric or abstract. Note that the face relation defines a canonical partial order on K. We call a simplex P a *facet* of another simplex

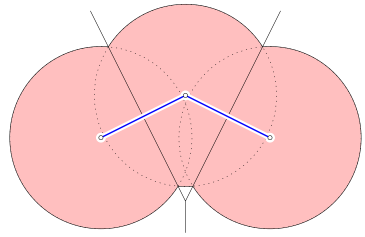

Figure 2: The Delaunay complex for the given radius has three vertices and two edges. In contrast, the Delaunay-Čech complex (not shown) has three vertices, three edges, and a triangle.

Q iff P is a face of Q and $\dim P = \dim Q - 1$, and we call (P, Q) a facet pair. The *Hasse diagram* of K is the directed graph whose nodes are the simplices and whose arcs are the facet pairs. It is the transitive reduction of the face relation. A *discrete vector field* is a partition V of the faces of K into singleton sets and facet pairs.

The Hasse diagram of K induces a digraph on V, denoted as $\mathcal{H}(V)$, obtained by contracting each facet pair in V. Specifically, the nodes of $\mathcal{H}(V)$ are the singletons and pairs in V, and there is an arc from μ to ν if there are simplices $P \in \mu$ and $Q \in \nu$ such that (P, Q) is a facet pair. The discrete vector field is *acyclic*, or a *discrete gradient field*, if $\mathcal{H}(V)$ is acyclic. In this case, the transitive closure of $\mathcal{H}(V)$ is a partial order on V, and we call $\mathcal{H}(V)$ the *Hasse diagram* of V. To prove acyclicity, we may construct a function $f\colon K \to \mathbb{R}$ that satisfies $f(P) = f(Q)$ whenever (P, Q) is a facet pair in V, and $f(P) < f(Q)$ whenever (P, Q) is a facet pair not in V. We call such an f a *discrete Morse function* that has V as its gradient. A simplex that does not belong to a pair in V is a *critical simplex*, and the corresponding value is a *critical value* of f.

To provide an intuition for the concept, we note that the pairs in a discrete gradient field correspond to elementary collapses (see e.g. [15, Chapter III]), except that the lower-dimensional simplex does not have to be free. An elementary collapse can be realized continuously by a deformation retraction. This implies that if we can transform a simplicial complex K to another simplicial complex K_0 using a sequence of such collapses, denoted by $K \searrow K_0$, then the two complexes have the same homotopy type. Indeed, the relation is slightly stronger, which is usually expressed by saying that K and K_0 are *simple-homotopy equivalent* [11]. A discrete Morse function can encode a simplicial collapse [17]:

Theorem 3.1 (Collapsing Theorem)**.** *If $f : K \to \mathbb{R}$ is a discrete Morse function and all critical simplices of f are contained in a subcomplex $K_t = f^{-1}(-\infty, t]$, then $K \searrow K_t$.*

Generalized discrete Morse theory To generalize the concepts of discrete Morse theory, we recall that an *interval* of the poset K, with the partial order given by the face relation, is a subset of the form

$$[P, R] = \{Q \mid P \subseteq Q \subseteq R\}.$$

The interval is non-empty iff P is a face of Q. In this case, the interval contains both simplices, and we refer to P as the *lower bound* and to R as the *upper bound* of the interval. Next, consider a partition of K into intervals. Borrowing from the nomenclature of [19], we call this set of intervals a *generalized discrete vector field*, denoting it by V. Indeed, a discrete vector field is the special case in which all intervals are either singletons or pairs. The corresponding directed graph, denoted again as $\mathcal{H}(V)$, has the intervals as its nodes and an arc from μ to ν if there are simplices $P \in \mu$ and $Q \in \nu$ such that (P, Q) is a facet pair in K. The generalized discrete vector field is *acyclic* if $\mathcal{H}(V)$ is acyclic, and if it is, then we call V a *generalized discrete gradient field* and $\mathcal{H}(V)$ the *Hasse diagram* of V. As before, acyclicity of V is asserted by the existence of a function $f\colon K \to \mathbb{R}$ that satisfies $f(P) = f(Q)$ whenever P is a face of Q and both belong to the same interval, while $f(P) < f(Q)$ whenever P is a face of Q but they belong to different intervals. We call such an f a *generalized discrete Morse function* that has V as its (generalized) gradient. We call an interval of cardinality 1 *singular* and the simplex it contains *critical*. Correspondingly, the value of a critical simplex is a *critical value* of f.

It is easy to see that for every generalized discrete gradient field, there is a discrete gradient field that refines the non-singular intervals into pairs. Both have the same critical simplices, implying that 3.1 also applies to generalized Morse functions. The refinement is in general not unique.

3.2 Radius Functions

The proximity complexes introduced in 3.1 have equivalent dual definitions, which are more convenient for our purposes. We begin by introducing these dual definitions, together with a precise statement of the general position assumption.

Čech and Delaunay functions A *circumsphere* of a finite set $Q \subseteq \mathbb{R}^n$ is an $(n-1)$-sphere that passes through all points of Q. A sufficient condition for the existence of a circumsphere is that Q is affinely independent. For $p < n$, the circumsphere is not unique, but there is only one *smallest circumsphere*, namely the one whose center lies in the affine hull of Q. We formulate conditions under which the smallest circumspheres do not contain more than the necessary points.

Definition 3.2 (General position)**.** *A finite set $X \subseteq \mathbb{R}^n$ is in* general position *if for every subset Q of at most $n + 1$ points in X,*

1. *Q is affinely independent, and*
2. *the smallest circumsphere of Q does not pass through any points of $X \backslash Q$.*

From now on, let X be a finite set of points in general position in $\mathbb{R}^n$. A circumsphere of a subset Q is *empty* if the open ball bounded by the sphere contains no point of X. It is not difficult to see that the Delaunay triangulation, as defined in 3.1, consists of all simplices $Q \subseteq X$ that have empty circumspheres; see also [12]. By Condition (i) of our general position assumption, all sets Q that have empty circumspheres are affinely independent, which suffices for $D(X)$ to be geometrically realizable in $\mathbb{R}^n$.

To make the step from the Delaunay triangulation to the Delaunay complexes for a radius, we note that among the empty circumspheres of a set $Q \in D(X)$, there is a unique one with smallest radius. We denote this *smallest empty circumsphere* by $S_D(Q)$. In addition, we introduce the *Delaunay function*, $f_D\colon D(X) \to \mathbb{R}$, defined by mapping Q to the squared radius of $S_D(Q)$. We will see shortly that f_D can be used to characterize Delaunay complexes.

Before we get there, we note that an *enclosing sphere* of $Q \subseteq \mathbb{R}^n$ is an $(n-1)$-sphere S such that the closed ball $\operatorname{conv} S$ contains all points of Q. Among all enclosing spheres, there is a unique one with smallest radius. We denote this *smallest enclosing sphere* by $S_C(Q)$, and we introduce the *Čech function*, $f_C\colon \Delta(X) \to \mathbb{R}$, defined by mapping Q to the squared radius of $S_C(Q)$. With these two functions, it is easy to say which simplices belong to the complexes.

Lemma 3.3 (Radius Functions Lemma)**.** *Let X be a finite set of points in general position in $\mathbb{R}^n$. Then*

$$C_r(X) = {f_C}^{-1}[0, r^2], \tag{5}$$
$$D_r(X) = {f_D}^{-1}[0, r^2], \tag{6}$$

for every $r \geqslant 0$.

Proof. A subset Q of X belongs to the Čech complex for radius r iff the closed balls with radius r centered at the points in Q have a non-empty common intersection. The points in this intersection are precisely the centers of all enclosing spheres of radius r. This implies (5).

The subset Q belongs to the Delaunay complex for radius r iff the Voronoi cells of the points in Q have a common intersection that includes points at a distance at most r from the points in Q. These points are precisely the centers of all empty circumspheres of radius at most r. This implies (6). □

It is easy to see that $f_C(Q) \leqslant f_D(Q)$ for every $Q \in D(X)$. This gives an alternative proof for the fact that $D_r(X) \subseteq DC_r(X) \subseteq C_r(X)$ for every $r \geqslant 0$.

Čech intervals We will prove collapsibility using structural properties of the two radius functions. We begin with the easier case of the Čech function. Given a simplex $Q \subseteq X$, recall that $S_C(Q)$ is the smallest enclosing sphere of Q. We find two simplices bounding an interval that consists of the simplices with the same smallest enclosing sphere as Q:

$$L_C(Q) = Q \cap S_C(Q);$$
$$U_C(Q) = X \cap \operatorname{conv} S_C(Q).$$

Note that $L_C(Q)$ is the unique simplex such that $S_C(Q)$ is the smallest circumsphere of $L_C(Q)$. Write V_C for the collection of such intervals $[L_C(Q), U_C(Q)]$, over all $Q \subseteq X$. We show that V_C is a generalized discrete gradient field.

Lemma 3.4 (Čech Function Lemma)**.** *Let X be a finite set of points in general position in $\mathbb{R}^n$. Then $f_C\colon \Delta(X) \to \mathbb{R}$ is a generalized discrete Morse function with gradient V_C.*

Proof. Let $Q \subseteq X$ and write $L = L_C(Q)$, $U = U_C(Q)$, and $S = S_C(Q)$. Clearly $L \subseteq Q \subseteq U$. Furthermore, every other simplex Q' with $L \subseteq Q' \subseteq U$ defines the same smallest enclosing sphere and therefore also the same simplices $L_C(Q') = L$ and $U_C(Q') = U$. Hence, the interval $[L, U]$ consists exactly of the simplices Q' with the same smallest enclosing sphere $S_C(Q') = S$.

Note that V_C is a generalized discrete vector field: its intervals partition $\Delta(X)$ since $S_C(Q') = S_C(Q)$ whenever $L_C(Q) \subseteq Q' \subseteq U_C(Q)$. It remains to show that V_C is a generalized gradient. But this is clear because V_C is defined by f_C, which is non-decreasing along increasing chains of the face relation, and it stays constant only within the intervals of V_C. □

We recall that our general position assumption implies that a simplex $Q \subseteq X$ has a smallest circumsphere iff $\dim Q \leqslant n$, and that the smallest circumspheres of different simplices are different. We thus have $\sum_{p=0}^{n} \binom{\dim X+1}{p+1}$ smallest circumspheres and the same number of intervals. Very few of them are singular, namely only the ones for which $L_C(Q) = U_C(Q)$. In other words, the critical simplices are the $Q \subseteq X$ such that the smallest enclosing sphere of Q is also an empty circumsphere. All these simplices also belong to the Delaunay triangulation.

Delaunay intervals Given a simplex $Q \in D(X)$, we now find two simplices that enclose between them the simplices with the same smallest empty circumsphere. To motivate the definition, we first investigate the intersections of simplices having a common smallest empty circumsphere:

Lemma 3.5 (Delaunay Face Lemma)**.** *Let X be a finite set of points in general position in $\mathbb{R}^n$. If $Q, Q' \subseteq X$ have the same smallest empty circumsphere, then this is also the smallest empty circumsphere of $Q \cap Q'$.*

Proof. Setting $P = Q \cap Q'$, we note that $\mathrm{Vor}\,(P)$ contains both $\mathrm{Vor}\,(Q)$ and $\mathrm{Vor}\,(Q')$ as faces. Indeed, $\mathrm{Vor}\,(P)$ is the smallest common intersection of Voronoi cells with this property. The center of the common smallest empty circumsphere satisfies $z \in \mathrm{Vor}\,(Q) \cap \mathrm{Vor}\,(Q')$ and therefore $z \in \mathrm{Vor}\,(P)$. Let now x be a point in P. Among the points in $\mathrm{Vor}\,(Q)$, z minimizes the distance to every point in Q and therefore also to x. Equivalently, the $(n-1)$-plane H normal to $x - z$ that passes through z *weakly separates* x from $\mathrm{Vor}\,(Q)$, i.e., z is contained in one closed half-space bounded by H, and $\mathrm{Vor}\,(Q)$ is contained in the other; see [7, Section 4.2.3]. Similarly, H weakly separates x from $\mathrm{Vor}\,(Q')$. Since $\mathrm{Vor}\,(P)$ is the smallest common intersection of Voronoi cells that contains $\mathrm{Vor}\,(Q)$ and $\mathrm{Vor}\,(Q')$ as faces, this implies that H also weakly separates z from $\mathrm{Vor}\,(P)$, and hence, z minimizes the distance to x among all points in $\mathrm{Vor}\,(P)$. We conclude that P has the same smallest empty circumsphere as Q and Q'. □

Writing $S_D(Q)$ for the smallest empty circumsphere of Q, the lower bound, $L_D(Q)$, is the intersection of all faces of Q whose smallest empty circumsphere is $S_D(Q)$, and the upper bound, $U_D(Q)$, consists of all points on $S_D(Q)$:

$$L_D(Q) = \bigcap_{S_D(Q') = S_D(Q)} Q';$$
$$U_D(Q) = X \cap S_D(Q').$$

Note that $U_D(Q)$ is the unique simplex such that $S_D(Q)$ is the smallest circumsphere of $U_D(Q)$. Write V_D for the collection of intervals $[L_D(Q), U_D(Q)]$, over all simplices $Q \in D(X)$. Using 3.5, we show that V_D is a generalized discrete gradient field.

Lemma 3.6 (Delaunay Function Lemma)**.** *Let X be a finite set of points in general position in $\mathbb{R}^n$. Then $f_D \colon D(X) \to \mathbb{R}$ is a generalized discrete Morse function with gradient V_D.*

Proof. Given $Q \in D(X)$, write $P = L_D(Q)$, $R = U_D(Q)$, and $S = S_D(Q)$. We first show that $S_D(Q') = S$ iff $P \subseteq Q' \subseteq R$. By definition, R contains all points of X on S, so $S_D(Q') = S$ implies $Q' \subseteq R$. Hence, it remains to show that S is the smallest empty circumsphere of $Q' \subseteq R$ iff $P \subseteq Q'$. If $P \nsubseteq Q'$ then, by definition of P, the simplex Q' has an empty circumsphere that is smaller than S. On the other hand, if $P \subseteq Q'$, then Q' does not have an empty circumsphere smaller than S, since any empty circumsphere of Q' is also an empty circumsphere of $P \subseteq Q'$, and S is the smallest empty circumsphere of P according to 3.5.

The intervals of V_D are of the form $[L_D(Q), U_D(Q)]$, over all $Q \in D(X)$. These intervals partition the Delaunay triangulation. Finally, f_D is non-decreasing along increasing chains of the face relation, and stays constant only within the constructed intervals, so V_D is the gradient of f_D. □

We have $L_D(Q) = U_D(Q)$ iff the center of $S_D(Q)$ is contained in $\mathrm{conv}\, U_D(Q)$, and if it is, then our general position assumption implies that it is an interior point of the simplex. This case is noteworthy since it implies that every singular interval of f_D is also a singular interval of f_C. Conversely, every singular interval of f_C is a singular interval of f_D. In other words, f_C and f_D have the same critical simplices. We summarize:

Lemma 3.7 (Critical Value Lemma)**.** *A simplex Q is a critical simplex of V_D iff it is a critical simplex of V_C iff $f_D(Q) = f_C(Q)$.*

Radon's theorem We prepare the analysis of the relation between Čech and Delaunay intervals by discussing a variant of Radon's theorem. In the original form, it asserts that any $p+1$ points in $\mathbb{R}^{p-1}$ can be partitioned into two non-empty sets whose convex hulls have a non-empty common intersection [24]. We are interested in a version of this theorem in which the $p+1$ points are the vertices of a p-simplex. Writing $R = \{x_0, x_1, \ldots, x_p\}$, we express the center of the smallest circumsphere as an affine combination:

$$z = \zeta_0 x_0 + \zeta_1 x_1 + \ldots + \zeta_p x_p.$$

The coefficients satisfy $\zeta_0 + \zeta_1 + \ldots + \zeta_p = 1$, but they are not necessarily all non-negative because z is not necessarily contained in $\operatorname{conv} R$. The p-simplex has $p+1$ *facets*, namely $R_i = R \backslash \{x_i\}$ for $0 \leqslant i \leqslant p$. We call R_i a *front facet* if $\zeta_i < 0$ and a *back facet* if $\zeta_i > 0$. Assuming general position, z does not lie in the $(p-1)$-plane of any facet, so every facet is either front or back. It is convenient to re-index such that R_0 to R_{j-1} are front facets and R_j to R_p are back facets. We call

$$F = \{x_j, x_{j+1}, \ldots, x_p\},$$
$$B = \{x_0, x_1, \ldots, x_{j-1}\}.$$

the *smallest front face* and the *smallest back face* of R, noting that F is either R or the common intersection of all front facets, and B is the possibly empty common intersection of all back facets. Observe that each front facet has the same smallest empty circumsphere as R, while every back facet has a smaller such sphere. If $R = U_D(Q)$ is an upper bound of a Delaunay interval, 3.5 asserts that $F = L_D(R)$ is the smallest face that shares the smallest empty circumsphere with R. The new version of Radon's theorem is an assertion about the position of the circumcenter relative to these two faces.

Lemma 3.8 (Front-Back Separation Lemma)**.** *Let R be a set of $p+1$ points in general position such that the center z of the smallest circumsphere lies outside* $\operatorname{conv} R$. *Any $(p-1)$-plane that weakly separates the smallest front face from the smallest back face also weakly separates z from the smallest back face.*

Proof. Assume without loss of generality that the points of R lie in $\mathbb{R}^p$ and the center of the smallest circumsphere lies at the origin. Writing this center as the affine combination of the $p+1$ vertices of R, we get $0 = \zeta_0 x_0 + \zeta_1 x_1 + \ldots + \zeta_p x_p$. By assumption of general position, the origin does not belong to the $(p-1)$-plane defined by any facet, which implies that all coefficients are non-zero. We re-index such that ζ_0 to ζ_{j-1} are negative and ζ_j to ζ_p are positive, define

$$y_F' = \zeta_j x_j + \zeta_{j+1} x_{j+1} + \ldots + \zeta_p x_p,$$
$$y_B' = \zeta_0 x_0 + \zeta_1 x_1 + \ldots + \zeta_{j-1} x_{j-1},$$

and observe that $y_F' = -y_B'$. Scaling the two points to $y_F = y_F'/(\zeta_j + \zeta_{j+1} + \ldots + \zeta_p)$ and $y_B = y_B'/(\zeta_0 + \zeta_1 + \ldots + \zeta_{j-1})$, we note that they are convex combinations of F and B, and that

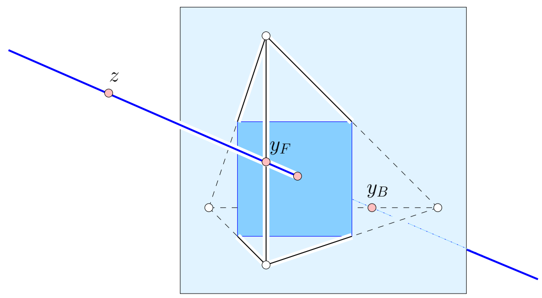

Figure 3: The line that passes through the circumcenter and intersects the smallest front as well as the smallest back face, and a plane that separates these two faces.

the two points are still collinear with the circumcenter at the origin; see 3. By construction, the order of the points along the common line is such that y_F lies between z and y_B. Any $(p-1)$-plane that weakly separates F from B intersects this line in a point between y_F and y_B. It then follows that this $(p-1)$-plane also weakly separates z from y_B, as claimed. □

To elucidate the connection to the Radon Theorem, we note that the central projections of F and of B from the circumcenter to a suitable $(p-1)$-plane give the Radon partition of the projection of R.

Interval intersection By 3.7, the Čech and the Delaunay functions have the same critical simplices, which form singular intervals. All other intervals are non-singular. While Delaunay intervals tend to be smaller than Čech intervals, there is no particular relation between them. Consider for example $L = L_C(Q)$ and $P = L_D(Q)$, for some $Q \in D(X)$. As illustrated in 4, it is possible that

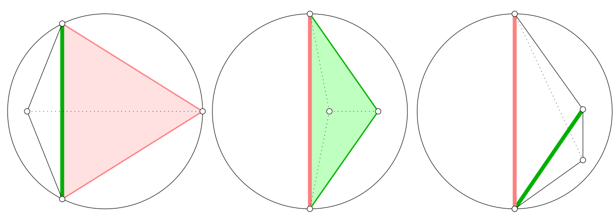

Figure 4: Three examples of a Delaunay tetrahedron within its smallest enclosing sphere, viewed from the center of the circumsphere. The circle is the intersection of the two spheres. From left to right: the lower bound of the Čech interval is a (pink) triangle, edge, edge, and the lower bound of the Delaunay interval is a (green) edge, triangle, edge.

P is a proper subset of L, that L is a proper subset of P, and that L and P are incomparable. Nevertheless, we have one important relationship, namely that the intersection of two non-singular intervals is non-singular.

Lemma 3.9 (Excluded Singularity Lemma)**.** *Let X be a finite set of points in general position in $\mathbb{R}^n$. The intersection of a non-singular Čech interval and a non-singular Delaunay interval is a possibly empty non-singular interval.*

Proof. Writing $[L, U]$ for the Čech interval and $[P, R]$ for the Delaunay interval, we assume that their intersection is non-empty. This common intersection is necessarily an interval, with lower bound $L \cup P$ and upper bound $U \cap R$. Setting $Q = L \cup P$, we aim to show that there exists a point $x \in U \cap R$ that does not belong to Q. It then follows that $Q \cup \{x\}$ is a second simplex in the intersection, implying that the intersection is not singular, as desired.

To establish the existence of such a point x, we set $S_0 = S_C(Q)$ and $S_1 = S_D(Q)$. Let H be the $(n-1)$-plane that contains $S_0 \cap S_1$, which is an $(n-2)$-sphere. We orient H such that S_1 encloses S_0 in the open half-space *in front* of H, and S_0 encloses S_1 in the open half-space *behind* H; see Figure 5. The center z_0 of S_0 is interior to $\operatorname{conv} L$, and L is a face of R contained in H. Because R

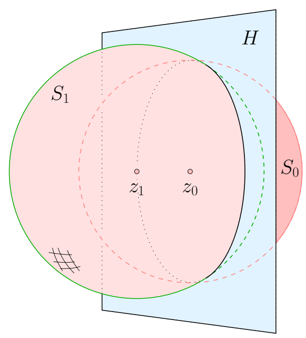

Figure 5: The plane contains the center of the smallest enclosing sphere and the circle in which this sphere intersects the smallest empty circumsphere.

is non-critical, by assumption, the center z_1 of S_1 lies outside $\operatorname{conv} R$ and in front of H. All points of R lie on S_1, and the points of $Q \subseteq R$ lie on or behind H. By definition of the Delaunay lower bound, $P = F$ is the smallest front face of R. Recall that $B = R \backslash F$ is the smallest back face of R. If $Q = R$, then $Q = F \cup L$ and $R = F \cup B$ imply $B \subseteq L$. Hence, B lies on H, $F \subseteq Q$ lies on or behind H, and H weakly separates F from B. By 3.8, F lies on or in front of H, implying that all points of R lie on H and therefore on $S_0 \cap S_1$. In particular, S_0 is a circumsphere of R. But S_0 has a smaller radius than S_1, contradicting the fact that S_1 is the smallest circumsphere of R. We conclude that Q does not contain all points of R.

We aim at proving that at least one of the points in $R \backslash Q$ lies behind H. If so, then this point x is enclosed by S_0 and therefore $x \in U \cap R$, as desired. To derive a contradiction, we assume that all points of $R \backslash Q$ lie in front of H. All points of L lie on H, so all points of B lie on or in front of H. Hence, H weakly separates F from B. But both $R \backslash Q \subseteq B$ and z_1 are in front of H, so H does not weakly separate z_1 from B, which contradicts 3.8. □

3.3 Collapsing Hierarchy

In this section, we prove our main result, which is the collapsibility of the Delaunay-Čech to the Delaunay complex, and of the Delaunay complex to the wrap complex. We begin with a sufficient condition for collapsibility.

Compatible restrictions We prove our collapsibility results using a structural insight into gradients. Let V be a generalized discrete gradient field on a simplicial complex K, let $K_0 \subseteq K$ be a subcomplex, and let V_0 be the restriction of V to K_0. We say that V *restricts compatibly* to K_0 if V_0 is again a generalized discrete gradient field, and a singular interval belongs to V iff it belongs to V_0.

Lemma 3.10 (Compatible Restriction Lemma)**.** *Let $K_0 \subseteq K$ be a pair of simplicial complexes. If there exists a generalized gradient V on K that restricts compatibly to K_0, then $K \searrow K_0$.*

Proof. We prove the claim by refining V to a (non-generalized) discrete gradient field W on K that restricts compatibly to K_0. Let $[L, U]$ be a non-singular interval in V. Its intersection with K_0 is either empty or a non-singular interval $[L, T]$ in K_0. In the first case, we choose an arbitrary vertex $x \in U \backslash L$ and decompose $[L, U]$ into pairs $[Q, Q \cup \{x\}]$ for all $x \notin Q \in [L, U]$. The second case is similar, except we make sure to choose $x \in T \backslash L$. Clearly, W is a discrete vector field on K. We note that $Q \in K_0$ iff $R \in K_0$ for every pair $[Q, R] \in W$. Since V restricts compatibly to K_0, this implies that W restricts compatibly to K_0. To prove that W is acyclic, we construct a discrete Morse function, $g\colon K \to \mathbb{R}$, with gradient W from a generalized discrete Morse function, $f\colon K \to \mathbb{R}$, with generalized gradient V. Specifically, for each singular $[Q, Q] \in W$, we set $g(Q) = f(Q)$, and for each pair $[Q, R] \in W$ contained in $[L, U] \in V$, we set

$$g(Q) = g(R) = \begin{cases} f(L) + \varepsilon \dim R & \text{if } Q, R \in K_0, \\ f(L) + \varepsilon \dim R + \lambda & \text{if } Q, R \notin K_0, \end{cases}$$

where $\varepsilon > 0$ is less than the smallest absolute difference between different values of f divided by the maximum dimension of any simplex, and λ is greater than the largest difference between values of f. It is straightforward to verify that g is a discrete Morse function with gradient W. Finally, to prove that K collapses to K_0, we let $m = \max_{Q \in K} f(Q)$ and $m_0 = \max_{Q \in K_0} f(Q)$ so that $K = g^{-1}(-\infty, m]$ and $K_0 = g^{-1}(-\infty, m_0]$. Moreover, since all critical simplices of V belong to K_0, no critical value of g is larger than m. Now 3.1 yields $K \searrow K_0$. ☐

Wrap complex Recall that the Delaunay function defines a generalized discrete gradient field, represented by the Hasse diagram, $\mathcal{H}(V_D)$, whose nodes are the Delaunay intervals. Every critical simplex forms a singular interval $[Q, Q]$ and therefore its own node in $\mathcal{H}(V_D)$. The *lower set* of a critical simplex $Q \in D(X)$, denoted by $\downarrow Q$, is the collection of simplices contained in intervals from which this node can be reached along directed paths. This concept is akin to the stable manifold of a critical point in smooth Morse theory, except that the lower sets of the critical simplices do not necessarily form a partition. Indeed, the lower sets can overlap, and some of the simplices may not belong to any lower set. The latter can be considered to belong to the lower set of the "outside", but it will not be necessary to formalize this intuition. The *wrap complex* for $r \geqslant 0$ is the union of the lower sets of all critical simplices Q with smallest empty circumsphere of radius at most r:

$$W_r(X) = \bigcup_{\substack{[Q,Q] \in V_D \\ f_D(Q) \leqslant r^2}} \downarrow Q.$$

Clearly, $W_r(X) \subseteq W_s(X)$ whenever $r \leqslant s$. Alternatively, we can define the wrap complexes as sublevel sets of another function, namely of $f_W\colon D(X) \to \overline{\mathbb{R}}$ defined by mapping $P \in D(X)$ to the minimum $f_D(Q)$ of any critical simplex Q for which $P \in \downarrow Q$, and to ∞ if P does not belong to the lower set of any critical simplex. Note that $f_D(P) \leqslant f_W(P)$ for every P, which implies $W_r(X) \subseteq D_r(X)$ for every $r \geqslant 0$.

Collapses We are ready to state and prove the main result.

Theorem 3.11 (Main Theorem)**.** *Let X be a finite set of points in general position in $\mathbb{R}^n$. Then*

$$DC_r(X) \searrow D_r(X) \searrow W_r(X),$$

for every $r \geqslant 0$.

Proof. We use 3.10 for both collapsibility results. Consider first the case in which $K = DC_r(X)$ and $K_0 = D_r(X)$. The generalized discrete gradient field is obtained by intersecting the Delaunay intervals with the Čech intervals:

$$V = \{[L_C(Q), U_C(Q)] \cap [L_D(Q), U_D(Q)] \mid Q \in D(X)\}.$$

To see that V is indeed acyclic, we consider $f\colon D(X) \to \mathbb{R}$ defined by $f(Q) = \frac{1}{2}(f_C(Q) + f_D(Q))$. It is not difficult to see that f has the described partition V of the Delaunay triangulation as its gradient. Each interval of V is either disjoint of K or contained in K, which implies that the restriction of V to K is the gradient of the restriction of f to K. 3.7 implies that all critical simplices of K are contained in K_0. Together with 3.9, this implies that the restriction of V to K restricts compatibly to K_0. Thus, by 3.10, we have $DC_r(X) \searrow D_r(X)$.

Consider second the case in which $K = D_r(X)$ and $K_0 = W_r(X)$. Here we use the Delaunay intervals V_D. Each Delaunay interval is either disjoint of K or contained in K, and similar for K_0. Moreover, by construction, a critical simplex belongs to K iff it belongs to K_0. Hence, the generalized discrete gradient field compatibly restricts to K_0, and 3.10 implies $D_r(X) \searrow W_r(X)$. □

3.4 Weighted Points

The weighted case arises when we use balls of different size in the construction of a complex. Encoding the difference by giving weights to the points, we use a real-valued parameter to grow and shrink all balls simultaneously. This concept is well known for Voronoi diagrams and Delaunay triangulations, whose weighted versions are sometimes referred to as *power diagrams* [5] and *regular triangulations* [20].

Weighted complexes It is not difficult to extend the four types of complexes considered to the weighted case. Let X be a finite set of points in $\mathbb{R}^n$, let $w\colon X \to \mathbb{R}$ be a *weight function*, and call $w(x)$ the *weight* of $x \in X$. The *weighted squared distance* of a point $a \in \mathbb{R}^n$ from x is the squared Euclidean distance minus the weight: $\pi_x(a) = \|x - a\|^2 - w(x)$. With this notion, we generalize the closed ball of radius r centered at x to $B_r^w(x) = \{a \in \mathbb{R}^n \mid \pi_x(a) \leqslant r^2\}$. Its center is x and its squared radius is $w(x) + r^2$, which can also be negative, in which case the ball is the empty set. Indeed, we allow the parameter to be negative, denoting it by $r^2 \in \mathbb{R}$ to remind us of its connection to the squared radius in the unweighted case. Technically, we let r take on values in $\sqrt{\mathbb{R}}$, which we define as the non-negative real numbers together with the non-negative multiples of the imaginary unit. Note that $B_r^w(x) = B_r(x)$ if $w(x) = 0$. The *weighted Čech complex* for $r \in \sqrt{\mathbb{R}}$ is

$$C_r^w(X) = \{Q \subseteq X \mid \bigcap_{x \in Q} B_r^w(x) \neq \emptyset\},$$

which is isomorphic to the nerve of the collection of balls $B_r^w(x)$, $x \in X$. To define the other complexes, we still need the notion of the *weighted Voronoi cell* of $x \in X$, denoted by $\mathrm{Vor}^w(x)$, which is the set of points $a \in \mathbb{R}^n$ such that $\pi_x(a) \leqslant \pi_y(a)$ for all $y \in X$. Similarly to the unweighted case, the weighted Voronoi cells are convex polyhedra, any two polyhedra intersect at most along shared boundaries, and together the weighted Voronoi cells cover $\mathbb{R}^n$. There are also differences: the weighted Voronoi cell of x does not necessarily contain x, and it is possible that $\mathrm{Vor}^w(x)$ be empty. The *weighted Delaunay triangulation*, denoted by $D^w(X)$, is isomorphic to the nerve of the collection of weighted Voronoi cells. As in the unweighed case, we need a mild general position assumption for this to give a simplicial complex in $\mathbb{R}^n$. A precise statement of that assumption is given shortly. We can now define the *weighted Delaunay complex* and the *weighted Delaunay-Čech complex* for $r \in \sqrt{\mathbb{R}}$:

$$D_r^w(X) = \{Q \subseteq X \mid \bigcap_{x \in Q} [B_r^w(x) \cap \mathrm{Vor}^w(x)] \neq \emptyset\},$$
$$DC_r^w(X) = \{Q \in D^w(X) \mid \bigcap_{x \in Q} B_r^w(x) \neq \emptyset\}.$$

Clearly, $D_r^w(X) \subseteq DC_r^w(X) \subseteq C_r^w(X)$, for every $r \in \sqrt{\mathbb{R}}$, just like in the unweighted case.

Weighted radius functions The weighted complexes have equivalent dual definitions based on generalizations of circumspheres and enclosing spheres. Note that we can write the sphere with center $a \in \mathbb{R}^n$ and radius ϱ as the zero set of the weighted squared distance function from the point a with weight ϱ^2:

$$S_\varrho(a) = \{b \in \mathbb{R}^n \mid \|b - a\|^2 - \varrho^2 = 0\}.$$

We say a point $x \in \mathbb{R}^n$ with weight $w(x) \in \mathbb{R}$ is *orthogonal* to $S_\varrho(a)$ if $\|x - a\|^2 - \varrho^2 - w(x) = 0$. Similarly, we say x is *closer than orthogonal* if the left hand side of the equation is negative, and *further than orthogonal* if the left hand side is positive. Generalizing the notion of circumsphere, we call $S_\varrho(a)$ an *orthosphere* of Q if every weighted point in Q is orthogonal to the sphere. Given x with weight $w(x)$ and $a \in \mathbb{R}^n$, we can find a unique weight ϱ^2 such that x and a are orthogonal. In other words, a single weighted point has an n-dimensional family of orthogonal spheres. Generically, $p + 1$ weighted points have an $(n - p)$-dimensional family of orthogonal spheres. One of the two requirements in the definition of general position is that there be no families of dimension higher than $n - p$.

Definition 3.12 (Weighted general position)**.** *A finite set $X \subseteq \mathbb{R}^n$ with weight function $w\colon X \to \mathbb{R}$ is in* general position *if for every subset Q of at most $n + 1$ points in X,*

1. *the family of orthospheres of Q has dimension $n - p$, and*
2. *the smallest orthosphere of Q is not orthogonal to any weighted point in $X \backslash Q$.*

Using the squared radius as the measure of size, we generalizing the smallest empty circumsphere: the *Delaunay sphere* of $Q \in D^w(X)$ is the smallest $(n - 1)$-sphere, $S_D^w(Q)$, that is orthogonal to every weighted point in Q and orthogonal or further than orthogonal to every weighted point in $X \backslash Q$. The corresponding *weighted Delaunay function*, $f_D^w\colon D^w(X) \to \mathbb{R}$, is defined by mapping Q to the squared radius of $S_D^w(Q)$. Recalling that the squared radius is really the weight of the center of the sphere, we note that this squared radius may be negative. Generalizing the smallest enclosing sphere, we define the *Čech sphere* of $Q \subseteq X$ as the smallest $(n - 1)$-sphere, $S_C^w(Q)$, that is orthogonal or closer than orthogonal to every weighted point in Q. The corresponding *weighted*

Čech function, $f_C^w \colon \Delta(X) \to \mathbb{R}$, is defined by mapping Q to the squared radius of $S_C^w(Q)$. As for the weighted Delaunay function, the value of a simplex under the weighted Čech function may be negative. With these two functions, it is easy to characterize which simplices belong to the weighted Čech and weighted Delaunay complexes.

Lemma 3.13 (Weighted Radius Functions Lemma)**.** *Let X be a finite set of points with weight function $w \colon X \to \mathbb{R}$ in general position in $\mathbb{R}^n$. Then*

$$C_r^w(X) = (f_C^w)^{-1}(-\infty, r^2],$$
$$D_r^w(X) = (f_D^w)^{-1}(-\infty, r^2],$$

for every $r \in \sqrt{\mathbb{R}}$.

Intervals It is not difficult to establish that the two weighted radius functions are indeed generalized discrete Morse functions. We begin with the description of the lower and upper bounds of the intervals, which are straightforward generalizations of these bounds in the unweighted case. For V_C^w, the lower bound of the interval that contains $Q \subseteq X$ is the simplex $L = L_C^w(Q) \subseteq Q$ of weighted points orthogonal to the Čech sphere of Q, and the upper bound is the simplex $U = U_C^w(Q) \subseteq X$ of weighted points orthogonal or closer than orthogonal to this Čech sphere. For V_D^w, the upper bound of the interval that contains $Q \in D^w(X)$ is the simplex $R = U_D^w(Q)$ of weighted points orthogonal to the Delaunay sphere of Q. A facet R_i of R is a *front facet* if it has the same Delaunay sphere as R, and a *back facet*, otherwise. As in the unweighted case, the lower bound is the common intersection of the front facets, $P = L_D^w(Q)$. These lower and upper bounds define the intervals of the gradients of the two weighted radius functions.

Lemma 3.14 (Weighted Čech and Delaunay Functions Lemma)**.** *Let X be a finite set of points with weight function $w \colon X \to \mathbb{R}$ in general position in $\mathbb{R}^n$. Then $f_C^w \colon \Delta(X) \to \mathbb{R}$ and $f_D^w \colon D^w(X) \to \mathbb{R}$ are generalized discrete Morse functions with gradients V_C^w and V_D^w, respectively.*

A crucial step toward proving the collapsibility of the weighted Delaunay-Čech complex to the weighted Delaunay complex is the comparison of the two gradient fields. The situation is again analogous to the unweighted case. A simplex $Q \subseteq X$ is critical for f_C^w iff every weighted point in $X \backslash Q$ is further than orthogonal to the Čech sphere of Q. But then the Čech sphere is equal to the Delaunay sphere, and Q is critical for f_D^w. Conversely, if Q is critical for f_D^w then it is critical for f_C^w. Besides having the same critical simplices, we need that the non-singular intervals intersect in non-singular intervals.

Lemma 3.15 (Weighted Excluded Singularity Lemma)**.** *Let X be a finite set of points with weight function $w \colon X \to \mathbb{R}$ in general position in $\mathbb{R}^n$. The intersection of a non-singular Čech interval and a non-singular Delaunay interval is a possibly empty non-singular interval.*

The proof is omitted because it is almost verbatim the same as in the unweighted case; see 3.2.

Collapsibility in the weighted case Before generalizing 3.11 from the unweighted to the weighted case, we still need to introduce the *weighted wrap complex* for $r \in \sqrt{\mathbb{R}}$, which is the union of the lower sets of all critical simplices Q whose value under the weighted Delaunay function is at most r^2:

$$W_r^w(X) = \bigcup_{\substack{[Q,Q] \in V_D^w \\ f_D^w(Q) \leqslant r^2}} \downarrow Q.$$

Similar to the unweighted case, the lower sets do not necessarily partition the triangulation, and they do not even necessarily cover it.

Theorem 3.16 (Weighted Main Theorem)**.** *Let X be a finite set of points with weight function $w\colon X \to \mathbb{R}$ in general position in $\mathbb{R}^n$. Then*

$$DC_r^w(X) \searrow D_r^w(X) \searrow W_r^w(X),$$

for every $r \in \sqrt{\mathbb{R}}$.

The proof is omitted because it is almost verbatim the same as in the unweighted case presented in 3.3.

3.5 Discussion

This constitutes an important generalization, which not only opens the door to smaller complexes (see WP4 for an application using this result), but also allows us to consider more general reconstruction theorems. In particular, this work is being adapted to a better complex for studying the self-map (in addition to a more efficient algorithm). While the Delaunay is known to be difficult to construct in high dimension, these results along with similar techniques may be able to allow us to obtain better approximations of persistence efficiently.

References

[1] F.R. Gantmacher, *The Theory of Matrices*, Chelsea Publishing Company, New York, 1959, 374+276 p.

[2] B. Hartley, T.O. Hawkes, *Rings, Modules and Linear Algebra*, Chapman and Hall, London, 1970, 210 p.

[3] P. van Dooren, The Computation of Kronecker's Canonical Form of a Singular Pencil, Lin. Alg. Appl., 27 (1979), 103–140.

[4] H. Edelsbrunner, G. Jabłoński, M. Mrozek, The Persistent Homology of a Self-map, preprint 2013, 16 p.

[5] F. Aurenhammer. Power diagrams: Properties, algorithms and application. *SIAM Journal on Computing*, 16(1):78–96, 1987.

[6] K. Borsuk. On the imbedding of systems of compacta in simplicial complexes. *Fundamenta Mathematicae*, 35(1):217–234, 1948.

[7] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, England, 2004.

[8] K. Buchin, T. K. Dey, J. Giesen, and M. John. Recursive geometry of the flow complex and topology of the flow complex filtration. *Computational Geometry*, 40(2):115–137, 2008.

[9] F. Cazals, A. Parameswaran, and S. Pion. Robust construction of the three-dimensional flow complex. In *Proceedings 24th Annual Symposium on Computational Geometry*, pages 182–191, 2008.

[10] H. Chintakunta and H. Krim. Distributed boundary tracking using alpha and Delaunay-Čech shapes, 2013. Submitted for publication in special issue of *Discrete Applied Mathematics*. `arXiv:1302.3982`.

[11] M. M. Cohen. *A Course in Simple-Homotopy Theory*, volume 10 of *Graduate Texts in Mathematics*. Springer-Verlag, 1973.

[12] B. Delaunay. Sur la sphère vide. A la mémoire de Georges Voronoï. *Bulletin de l'Académie des Sciences de l'URSS*, (6):793–800, 1934.

[13] H. Edelsbrunner. The union of balls and its dual shape. *Discrete and Computational Geometry*, 13(1):415–440, 1995.

[14] H. Edelsbrunner. Surface reconstruction by wrapping finite sets in space. In B. Aronov, S. Basu, J. Pach, and M. Sharir, editors, *Discrete and Computational Geometry*, volume 25 of *Algorithms and Combinatorics*, pages 379–404. Springer-Verlag, 2003.

[15] H. Edelsbrunner and J. Harer. *Computational Topology. An Introduction.* American Mathematical Society, Providence, Rhode Island, 2010.

[16] H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Transactions on Graphics*, 13(1):43–72, 1994.

[17] R. Forman. Morse theory for cell complexes. *Advances in Mathematics*, 134(1):90–145, 1998.

[18] R. Forman. A user's guide to discrete Morse theory. *Séminaire Lotharingien de Combinatoire*, B48c:1–35, 2002.

[19] R. Freij. Equivariant discrete Morse theory. *Discrete Mathematics*, 309(12):3821–3829, 2009.

[20] I. M. Gelfand, M. M. Kapranov, and A. V. Zelevinsky. *Discriminants, Resultants and Multidimensional Determinants*. Mathematics: Theory & Applications. Birkhäuser, 1994.

[21] J. Giesen and M. John. The flow complex: A data structure for geometric modeling. *Computational Geometry*, 39(3):178–190, 2008.

[22] J. Giesen and L. Kuehne. A parallel algorithm for computing the flow complex. In *Proceedings 29th Annual Symposium on Computational Geometry*, pages 57–66, 2013.

[23] J. Leray. Sur la forme des espaces topologiques et sur les points fixes des représentations. *Journal de Mathématiques Pures et Appliquées*, 24:95–167, 1945.

[24] J. Radon. Mengen konvexer Körper, die einen gemeinsamen Punkt enthalten. *Mathematische Annalen*, 83(1-2):113–115, 1921.